



# Project Status Report

## High End Computing Capability Strategic Capabilities Assets Program

August 10, 2014


Dr. Rupak Biswas – Project Manager  
NASA Ames Research Center, Moffett Field, CA  
Rupak.Biswas@nasa.gov  
(650) 604-4411

# HECC Team Optimizes MITgcm Ocean Model to Scale to 70,000 Cores on Pleiades



- HECC visualization experts successfully optimized the I/O performance of the Massachusetts Institute of Technology (MIT) General Circulation Model (MITgcm) in a test run across 70,000 cores on the Pleiades supercomputer—doubling the scale of a landmark 35,000-core run in January 2014.
- HECC teams identified and resolved system-level issues with InfiniBand (IB) cabling, IB microcode, and the MPI communications library.
  - Re-implementation of MITgcm output routines allows results to be written to disk without slowing computation. Output for the previous timestep overlaps with computation of the next, allowing the simulation to proceed without waiting.
  - New I/O code can write at a sustained rate of over 10 gigabytes per second.
- With the success of this effort, researchers working on the Estimating the Circulation and Climate of the Ocean (ECCO) project are contemplating even higher resolution runs.

**Mission Impact:** High-resolution global ocean simulations are critical for researchers to understand how ocean, sea-ice, and atmospheric systems interact and evolve.



Dimitris Menemenlis, co-investigator on the ECCO project, with a printout of output from a high-resolution (1/48-degree) MITgcm simulation. The data was produced on 70,000 cores of Pleiades and rendered at the NASA Advanced Supercomputing (NAS) facility. Pixels were then sent over the network for live display at MIT.

**POCs:** Chris Henze, [chris.henze@nasa.gov](mailto:chris.henze@nasa.gov), (650) 604-3959, NASA Advanced Supercomputing Division; Bron Nelson, [bron.c.nelson@nasa.gov](mailto:bron.c.nelson@nasa.gov), (650) 604-4329, NASA Advanced Supercomputing Division, Computer Sciences Corp.

# HECC Engineers Upgrade Operating System (OS) on Lou Archive System to SLES11SP3



- HECC engineers upgraded the Lou archive system's OS to SUSE Linux Enterprise System 11 Service Pack 3 (SLES11SP3).
- The upgrade was timed to coincide with a planned Pleiades downtime, to minimize impact on users. During this time, the engineers also updated Lou's storage management software to InfiniteStorage Software Platform (ISSP) 3.1.
- The ISSP enhancements provide system administrators with more control over data migration, including the ability to cancel or re-order writes to the archive system.
- The upgrade was originally scheduled for May, but a performance issue delayed full deployment. HECC engineers worked closely with SGI to identify and develop a fix before completing the upgrade.

**Mission Impact:** Maintaining an up-to-date operating system and data management software helps ensure efficient operation of the Lou archive system, a critical resource for NASA users to store and access data produced on HECC supercomputers.



The Lou archive system is an SGI parallel Data Migration Facility cluster comprising four 32-processor front-end systems, each with 64 gigabytes of memory. The system stores two copies of over 36 petabytes of data at the NAS primary and secondary compute facilities.


**POCs:** Bob Ciotti, [bob.ciotti@nasa.gov](mailto:bob.ciotti@nasa.gov), (650) 604-4408, NASA Advanced Supercomputing (NAS) Division; Davin Chan, [davin.chan@nasa.gov](mailto:davin.chan@nasa.gov), (650) 604-3613, NAS Division, Computer Sciences Corp.

# InfiniBand Cable Replacement Improves Pleiades' Stability



- During a recent system maintenance window, SGI and Mellanox engineers replaced 917 Fourteen-Data-Rate (FDR) InfiniBand (IB) cables to improve the stability and reliability of the Pleiades IB fabric.
- Mellanox staff originally estimated that between 30% to 53.8% of the IB cables would need to be replaced due to a manufacturing defect.
- HECC engineers worked closely with SGI and Mellanox to identify the problematic cables, and determined that only 15.6% of the FDR cables needed to be replaced.
- This reduction in the number of replaced cables significantly shortened the time that Pleiades was unavailable to users.

**Mission Impact:** Ensuring a stable InfiniBand network is essential for providing reliable high-end computing resources to HECC scientific and engineering users.



This birds-eye-view of the rows of racks that comprise the Pleiades supercomputer shows some of the 65 miles of InfiniBand cabling that interconnects all the compute nodes.

**POCs:** Bob Ciotti, [bob.ciotti@nasa.gov](mailto:bob.ciotti@nasa.gov), (650) 604-4408, NASA Advanced Supercomputing (NAS) Division; Davin Chan, [davin.chan@nasa.gov](mailto:davin.chan@nasa.gov), (650) 604-3613, NAS Division, Computer Sciences Corp.



# Control Room Analysts Protect HECC Systems During Cooling Pump Failure



- In the early morning on Saturday, July 12, two of the four cooling pumps at the NASA Advanced Supercomputing (NAS) facility failed, putting at risk the facility's computing systems, including the Pleiades and Endeavour supercomputers.
- HECC Control Room analysts took immediate measures to safeguard the systems:
  - At 3:35 a.m. the analysts were notified of a problem with the chillers, and began monitoring floor temperatures in the computer room at several locations.
  - Within 15 minutes temperatures began to rise, and the analysts initiated procedures to shut down the compute nodes.
  - By 3:56 a.m., all systems had been shut down.
- This quick response prevented damage to the computing systems, averting a catastrophic failure and saving millions of dollars worth of hardware.
- Cooling pump repairs are now underway and Pleiades is operating at half capacity, while Endeavour has returned to full capacity.
- HECC teams are continuing to investigate the root cause of the pump failure.

**Mission Impact:** HECC's 24x7 Control Room support enabled a quick response to a potentially catastrophic incident, preventing millions of dollars of damage to NASA's flagship supercomputer, and reducing impact to users across all mission directorates.



Four cooling pumps, located in cooling towers outside the NAS facility, maintain a safe operating temperature for the facility's supercomputing systems. Two of these pumps failed during the facility incident on July 12.

**POCs:** Chris Buchanan, [chris.buchanan@nasa.gov](mailto:chris.buchanan@nasa.gov), (650) 604-4308,  
Leigh Ann Tanner, [leighann.tanner@nasa.gov](mailto:leighann.tanner@nasa.gov), (650) 604-4468,  
NASA Advanced Supercomputing Division, Computer Sciences Corp.

# New Version of Lumber Tool Used to Quickly Identify Critical System Issues



- HECC's Application Performance and Productivity team has extended its system log analysis tool, Lumber, to characterize the impact of different types of job failures.
- Lumber can now take as input a user-supplied shared object that has access to Lumber's internal data structures.
  - This feature allows a small piece of C code to be invoked as Lumber is processing log messages associated with PBS jobs.
  - With this extension, Lumber can efficiently identify message patterns that match known failure modes for jobs.
- HECC Control Room analysts monitor job failures reported by users and can easily add newly identified failure types to the Lumber extension. In addition, Lumber is now used to generate a Job Failure Summary for the weekly HECC technical report.

**Mission Impact:** Regular, timely reports on the impact of failed jobs identifies important system issues and helps HECC system administrators prioritize their work, leading to improved productivity for users.

## Job Failure Summary from:

Sat Jul 12 00:00:00 2014 to Fri Jul 18 23:59:59 2014

There were 8749 jobs in the time region,  
of which 296 indicate as failed.

Count	SBUs	Failure type
108	80254.44	The PBS Server discarded the job because it appeared a node was down
94	25094.16	>= 1 node associated with the job booted for unknown reasons
54	4541.09	>= 1 node associated with the job booted because of out of memory (oom)
53	37281.76	local QP error
22	4170.30	The PBS Server discarded the job for unknown reasons
12	9760.28	Job experienced out of memory
3	6835.66	Job deleted by request
1	833.45	>= 1 node associated with the job booted due to an uncorrectable ecc memory error

For each known failure mode, Lumber produces a job failure report that indicates the number of individual jobs experiencing that type of failure, as well as the total number of SBUs consumed by all those jobs.


**POCs:** David Barker, david.p.barker@nasa.gov, (650) 604-4292, NASA Advanced Supercomputing (NAS) Division, Supersmith; Chris Bell, christopher.m.bell@nasa.gov, (650) 604-4444, NAS Division, Computer Sciences Corp.

# Modeling Hurricanes and Other High-Impact Weather Systems in High Resolution \*



- The representation of clouds is a major challenge in climate modeling and prediction. Research scientists at Goddard Space Flight Center (GSFC) are running high-fidelity simulations on Pleiades to simulate high-impact weather systems—from severe thunderstorms to squall lines to hurricanes.
  - The simulations provide unique and detailed insights into the dynamics and precipitation processes associated with the formation of clouds, cloud systems, and hurricanes.
  - These simulations also provide 4D datasets to NASA's Tropical Rainfall Measuring Mission developers to improve the performance of their rainfall and latent heating algorithms.
  - In addition, the simulated results have been archived on a NASA cloud library website.
- Simulations were performed on HECC's Pleiades and GSFC's Discover systems using about 2,000 processors. Depending on the nature of the information required, the computational runs lasted from days to months, resulting in several petabytes of data.

**Mission Impact:** Simulations run on NASA's Pleiades supercomputer have provided unique and detailed insights into the processes associated with the formation of clouds, cloud systems, and hurricanes.



At bottom, an image from simulated cloud systems generated during the Tropical Warm Pool International Cloud Experiment. Blue indicates cold pool structures generated by the storm. Animations produced by the HECC visualization team clearly show the cold-pool induced convective cloud structures, as shown in the top image. Both simulations were run using the Goddard Cumulus Ensemble model at 1 km horizontal resolution. The results compare well with observed rainfall. *Tim Sandstrom, NASA/Ames*

**POCs:** Wei-Kuo Tao, [wei-kuo.tao-1@nasa.gov](mailto:wei-kuo.tao-1@nasa.gov), (301) 614-6269;  
Xiaowen Li, [xiaowen.li@nasa.gov](mailto:xiaowen.li@nasa.gov), (301) 614-6319; NASA  
Goddard Space Flight Center


\* HECC provided supercomputing resources and services in support of this work

# CFD Simulations Enable Verification of Wind Tunnel Results for Future Passenger Planes \*



- Aeronautics engineers at NASA and MIT are performing simulations on Pleiades to study the D8 “double-bubble” aircraft concept, which uses a wider fuselage, lightweight materials, and engines with boundary-layer ingesting nacelles mounted on the top rear of the fuselage.
  - These simulations are being performed in conjunction with wind tunnel tests at NASA Langley to determine whether the boundary-layer ingesting nacelles improve fuel efficiency.
  - Aerodynamic forces on the D8 are computed to determine the amount of engine power required in cruise flight. The power required is then compared to a traditional podded nacelle design to verify fuel savings.
  - Computational results are also expected to shed additional light on the reasons behind the better fuel efficiency by examining the airflow behavior in the vicinity of the engines.
- These CFD simulations use hundreds of Pleiades processors to compute the flow around the aircraft and the flow through a simulated engine fan.

**Mission Impact:** Computational verification of improved airplane designs, enabled by NASA’s Pleiades supercomputer, may ultimately lead to fuel savings and lower air travel costs for airline passengers.



Top: Visualization of the pressure field over a D8 with traditional podded and integrated nacelles. The pressure and viscous forces are integrated over the airplane surface to obtain the total force, and the fuel requirements of the two are compared. Scale models (bottom) are used in a wind tunnel to assess fuel savings. *Shishir Pandya, NASA/Ames*

**POC:** Shishir Pandya, [shishir.pandya@nasa.gov](mailto:shishir.pandya@nasa.gov), (650) 604-3981, NASA Ames Research Center

\* HECC provided supercomputing resources and services in support of this work



# HECC Facility Hosts Several Visitors and Tours in July 2014



- HECC hosted 3 tour groups in July; guests learned about the agency-wide missions being supported by Pleiades, and viewed the QuAIL system.
- NOTE: Due to facility outages and the hyperwall computer system upgrade, many tours have been postponed until the hyperwall is returned to production.
- Visitors this month included:
  - 30 students from the 2014 Ames Singularity University received an overview of the HECC project and NAS facility, and a quantum computer presentation and tour.
  - 32 high school students from the NASA Opportunities in Visualization, Art, and Science program, associated with the University of California, Berkeley and a NASA education grant, received a visualization presentation and computer room tour.
  - 15 undergraduate students from the Stanford University Army High Performance Computing Research Center summer program received an Ames tour, including the NAS facility.



High school students involved in a NASA/University of California, Berkeley education program were given a visualization presentation and a tour of the main NAS computer room, as part of their visit to Ames Research Center.

**POC:** Gina Morello, [gina.f.morello@nasa.gov](mailto:gina.f.morello@nasa.gov), (650) 604-4462, NASA Advanced Supercomputing Division

# Phi-based Systems



- **Background:**

*Two Xeon Phi-based systems are being utilized as pathfinding resources to determine whether the Many Integrated Core (MIC) Architecture is cost effective for NASA's computational requirements.*

- Maia is a 128-node SGI system with two Intel Xeon Phi accelerator cards in each node.
- Mira is a 64-node Cray system with two Intel Xeon Phi accelerator cards in each node.

- **Status**

- Both Maia and Mira are currently offline due to the cooling pump issue at the main NAS facility. The systems are expected to be back online in early August after the cooling pumps are fully restored.
- Two staff members from HECC, along with several members from other NASA centers, will attend an upcoming Intel “dungeon” (customer-focused lab) in Hillsboro, Oregon, to focus on application performance optimization for Xeon Phi platforms.

- **Upcoming Activities for August**

- Maia: The system should be available for testing early in the month.
- Mira: The “NASification” process will most likely continue for a few more weeks.

# Papers and Presentations



- **“NASA Advanced Computing Environment for Science & Engineering,”** R. Biswas, presented at the Ter@tec Forum, Palaiseau, France, July 1-2, 2014.
- **“Turbulence-Induced Relative Velocity of Dust Particles III: The Probability Distribution,”** L. Pan, P. Padoan, J. Scalo, arXiv:1407.1101 [astro-ph.EP], July 4, 2014. \*  
<http://arxiv.org/abs/1407.1101>
- **“Three-Point Galaxy-Galaxy Lensing as a Probe of Dark Matter Halo Shapes,”** S. Adhikari, C. Chue, N. Dalal, arXiv:1407.1128 [astro-ph.CO], July 4, 2014. \*  
<http://arxiv.org/abs/1407.1128>
- **“The Role of Subsurface Flows in Solar Surface Convection: Modeling the Spectrum of Supergranular and Larger Scale Flows,”** J. Lord, R. Cameron, M. Rast, M. Rempel, T. Roudier, arXiv:1407.2209 [astro-ph.SR], July 8, 2014. \*  
<http://arxiv.org/abs/1407.2209>
- **“Spectro-Polarimetric Properties of Small-Scale Plasma Eruptions Driven by Magnetic Vortex Tubes,”** I. Kitiashvili, arXiv:1407.2295 [astro-ph.SR], July 8, 2014. \*  
<http://arxiv.org/abs/1407.2295>
- **“Probing Satellite Quenching with Galaxy Clustering,”** R. Chamberlain, N. Dalal, A. Hearin, P. Ricker, arXiv:1407.2648 [astro-ph.GA], July 9, 2014. \*  
<http://arxiv.org/abs/1407.2648>

*\* HECC provided supercomputing resources and services in support of this work*

# Papers and Presentations (cont.)



- **“A Possible Approach for Big Data Access to Support Climate Science,”** M. Foster, presented at the CrossConnects workshop hosted by ESnet and Internet2, Boulder, CO, July 14-16, 2014. \*  
<http://meetings.internet2.edu/2014-cc-climate/detail/10003336/>
- **“Global Simulations of Axisymmetric Radiative Black Hole Accretion Disks in General Relativity with a Sub-grid Magnetic Dynamo,”** A. Sadowski, et al., arXiv: 1407.4421 [astro-ph.HE], July 16, 2014. \*  
<http://arxiv.org/abs/1407.4421>
- **“A Study of the Early-Stage Evolution of Relativistic Electron-Ion Shock Using Three-Dimensional Particle-in-Cell Simulations,”** E. Choi, K. Min, K.-I. Nishikawa, C. Choi, Physics of Plasmas (1994-Present), vol. 21, issue 7, July 17, 2014. \*  
<http://scitation.aip.org/content/aip/journal/pop/21/7/10.1063/1.4890479>
- **“Discovery of a Transiting Planet Near the Snow-Line,”** D. Kipping, et al., arXiv: 1407.4807 [astro-ph.EP], July 17, 2014. \*  
<http://arxiv.org/abs/1407.4807>
- **“Shock Corrugation by Rayleigh-Taylor Instability in Gamma-Ray Burst Afterglow Jets,”** P. Duffell, A. MacFadyen, The Astrophysical Journal Letters, vol. 791, no. 1, July 22 2014. \*  
<http://iopscience.iop.org/2041-8205/791/1/L1>

*\* HECC provided supercomputing resources and services in support of this work*



# Papers and Presentations (cont.)




- **“Inverse Cascade of Non-Helical Magnetic Turbulence in a Relativistic Fluid,”**  
J. Zrake, arXiv:1407.5626 [astro-ph.HE], July 21, 2014. \*  
<http://arxiv.org/abs/1407.5626>
- **AIAA/ASME/SAE/ASEE Joint Propulsion Conference**, Cleveland, OH, July 28-30, 2014.
  - **“Effect of Fuel Injection and Mixing Characteristics on Pulse-Combustor Performance at High-Pressure,”** S. Yungster, D. Paxson, H. Perkins. \*  
<http://arc.aiaa.org/doi/pdf/10.2514/6.2014-3728>
  - **“NASA Numerical and Experimental Evaluation of UTRC Low Emissions Injector,”** Y. Hicks, S. Tedder, R. Anderson, A. Iannetti, L. Smith, Z. Dai. \*  
<http://arc.aiaa.org/doi/pdf/10.2514/6.2014-3627>

*\* HECC provided supercomputing resources and services in support of this work*




- **Quantum Computing's Reality Fluctuates Over Time**, *SFGate*, July 3, 2014—Thomas Lee, Business Editor for the San Francisco Chronicle, discussed how quantum computing technology will change Silicon Valley, and spoke with HECC Project Manager Rupak Biswas about the agency's goals and how the D-Wave Two system at NASA Ames could be utilized.  
<http://www.sfgate.com/technology/article/Quantum-computing-s-reality-fluctuates-over-time-5599154.php>
- **Welcome to Quantum Computing**, *Baltimore Post-Examiner*, July 25, 2014—A detailed overview of quantum computing as it stands today, including a brief discussion on how NASA and Google have teamed up to test the D-Wave Two computer.  
<http://baltimorepostexaminer.com/welcome-quantum-computing/2014/07/25>

# HECC Utilization






July 2014

# HECC Utilization Normalized to 30-Day Month






# HECC Utilization Normalized to 30-Day Month




- ① Columbia 21, 23, and 24 retired, Endeavour 2 added
- ② Columbia 22 retired; Endeavour 1 added
- ③ 32 Harpertown Racks retired
- ④ 32 Harpertown Racks retired; 46 Ivy Bridge Racks added
- ⑤ 6 Ivy Bridge Racks added; 20 Nehalem and 12 Westmere Racks Retired
- ⑥ 8 Ivy Bridge Racks added mid-Feb; 8 additional Ivy Bridge Racks late Feb.
- ⑦ 4 Ivy Bridge Racks added mid-March

# Tape Archive Status




July 2014

# Tape Archive Status




# Pleiades: SBUs Reported, Normalized to 30-Day Month






# Pleiades: Devel Queue Utilization




# Pleiades: Monthly Utilization by Job Length




July 2014

# Pleiades: Monthly Utilization by Size and Mission




# Pleiades: Monthly Utilization by Size and Length




July 2014

# Pleiades: Average Time to Clear All Jobs




# Pleiades: Average Expansion Factor






# Endeavour: SBUs Reported, Normalized to 30-Day Month




# Endeavour: Monthly Utilization by Job Length




July 2014

# Endeavour: Monthly Utilization by Size and Mission




July 2014

# Endeavour: Monthly Utilization by Size and Length




July 2014

# Endeavour: Average Time to Clear All Jobs



# Endeavour: Average Expansion Factor





# Maia: SBUs Reported, Normalized to 30-Day Month

